What is claimed is:

1. A method for indexing a database table, said table comprising a column of values, the method comprising the steps of:

identifying a plurality of substrings, each substring comprising one or more characters; and

for each substring creating a corresponding index over the column, wherein each corresponding index includes a respective entry for each value in the column, the respective entry for a particular value being indicative of whether the corresponding substring occurs in that particular value.

2. The method according to claim 1, wherein the respective entry for the particular value is indicative of a number of occurrences of the corresponding substring in that particular value.

3. A method for running a query on a database table, said table comprising a column of values and including for each of a plurality of substrings, a corresponding index over the column, wherein each corresponding index includes a respective entry for each value in the column, the respective entry for a particular value being indicative of whether the corresponding substring occurs in that particular value, the method comprising the steps of:

identifying a set of values that potentially satisfy the query based on the plurality of corresponding indices.

4. The method according to claim 3, further comprising the steps of:

scanning the set of values to find matching values satisfying the query; and

generating a result set of the matching values.

5. The method according to claim 4, wherein the step of identifying the set of values includes the step of:

for each value, combining the respective entries from the corresponding indices.

6. The method according to claim 4, wherein the query includes a LIKE phrase.

7. The method according to claim 4, further comprising the steps of:

identifying a search term in the query;

separating the search term into a plurality of search substrings, each search substring corresponding to one of the plurality of substrings; and

for each substring corresponding to a search substring, building its corresponding index to generate a set of built indices.

8. The method according to claim 7, wherein each search substring has a length of two characters.

9. The method according to claim 7, wherein there are 100 search substrings representing, in character format, the numeric range from 00 to 99.

10. The method according to claim 7, wherein the step of identifying a set of values that potentially satisfy the query includes the step of:

for each value, combining its respective entries from each of the built indices to generate a composite entry, the composite entry indicative of whether all the search substrings occur within the value.

11. The method according to claim 10, wherein each respective entry is a binary value; and the respective entries are logically combined to generate the composite entry.

12. The method according to claim 10, further comprising the steps of:

building a bitmap over the column, each element of the bitmap corresponding to a value of the column and each bitmap element value determined based on the composite entry for the corresponding value, each such bitmap element value being on if the composite entry indicates that all the search substrings occur within the value, otherwise the value being off.

13. The method according to claim 12, wherein the step of scanning further includes:

scanning those values having a bitmap element with an on value.

14. The method according to claim 3, further comprising the steps of:

maintaining a history of received queries; and

creating one or more of the plurality of substrings based on the history.

15. The method according to claim 14, wherein the step of identifying one or more of the plurality of substrings includes the steps of:

tracking a number of occurrences of a particular search substring within the received queries; and

determining when the number of occurrences exceed a predetermined threshold.

16. An apparatus for executing a query on a database table, said table comprising a column of values, the apparatus comprising:

at least one processor;

a memory coupled with the at least one processor; and

a program code residing in the memory and executed by the at least one processor, the program code configured to:

create a corresponding index over the column, for each of a plurality of substrings, wherein the corresponding index includes a respective entry for each value in the column, the respective entry indicative of whether the substring occurs within the value, and

identify a set of values that potentially satisfy the query based on the plurality of corresponding indices.

17. An apparatus for executing a query on a database table, said table comprising a column of values, the apparatus comprising:

at least one processor;

a memory coupled with the at least one processor;

a plurality of indices stored within said memory each index corresponds to one of a plurality of substrings and each index includes a respective entry for each value in the column, the respective entry indicative of whether the corresponding substring occurs within the value; and

a program code residing in the memory and executed by the at least one processor, the program code configured to scan the values of the table based on a combination of the plurality of indices.

18. The method according to claim 17, wherein the program code is further configured to:

build a bitmap based on the plurality of indices, the bitmap having an element for each value of the column that is set based on the respective entries for that value from the plurality of indices.

19. A program product for indexing a database table, said table comprising a column of values, comprising:

program code configured upon execution thereof to:

identify a plurality of substrings, each substring comprising one or more characters, and

for each substring, create a corresponding index over the column, wherein each corresponding index includes a respective entry for each value in the column, the respective entry for a particular value being indicative of whether the corresponding substring occurs in that particular value; and

a signal bearing medium bearing the program code.

20. A program product for running a query on a database table, said table comprising a column of values and including for each of a plurality of substrings, a corresponding index over the column, wherein each corresponding index includes a respective entry for each value in the column, the respective entry for a particular value being indicative of whether the corresponding substring occurs in that particular value, the program product comprising:

program code configured upon execution thereof to:

identify a set of values that potentially satisfy the query based on the plurality of corresponding indices; and

a signal bearing medium bearing the program code.